

Diagnosis by Integrating Model-Based Reasoning with Knowledge-Based Reasoning

Tom Bylander
Laboratory for Artificial Intelligence Research
Department of Computer and Information Science
The Ohio State University
Columbus, Ohio 43210
byland@cis.ohio-state.edu

Abstract

Our research investigates how observations can be categorized by integrating a qualitative physical model with experiential knowledge. Our domain is diagnosis of pathologic gait in humans, in which the observations are the gait motions, muscle activity during gait, and physical exam data, and the diagnostic hypotheses are the potential muscle weaknesses, muscle mistimings, and joint restrictions. Patients with underlying neurological disorders typically have several malfunctions. Among the problems that need to be faced are: the ambiguity of the observations, the ambiguity of the qualitative physical model, correspondence of the observations and hypotheses to the qualitative physical model, the inherent uncertainty of experiential knowledge, and the combinatorics involved in forming composite hypotheses. Our system divides the work so that the knowledge-based reasoning suggests which hypotheses appear more likely than others, the qualitative physical model is used to determine which hypotheses explain which observations, and another process combines these functionalities to construct a composite hypothesis based on explanatory power and plausibility. We speculate that the reasoning architecture of our system is generally applicable to complex domains in which a less-than-perfect physical model and less-than-perfect experiential knowledge need to be combined to perform diagnosis.

The Promise of Deep Knowledge

Recently in knowledge-based systems research, there has been an emphasis on "deep" knowledge over "shallow" knowledge. Deep knowledge is based on the causal mechanisms underlying the domain, typically obtained through scientific studies and incorporated into physical models of the domain, while shallow knowledge is based on experiential knowledge, typically obtained from human experts and incorporated in rule-based systems [4]. The promise of deep knowledge is that the conclusions of a knowledge-based system be inferred from an accurate physical model of the domain, rather than dependent on a time-consuming and error-prone knowledge engineering effort.

The emphasis on deep knowledge has influenced research on diagnosis [9, 19]. In this line of research, the process of diagnosis is reduced to manipulating the physical

model within the knowledge-based system, according to the following presumptions.

1. Let M be the physical model which describes the domain when everything is functioning as it should.
2. Let \mathcal{M} be the set of all physical models consistent with the observations. If $M \notin \mathcal{M}$, then there is a malfunction.
3. If there is a malfunction, then for each $M' \in \mathcal{M}$, the differences between M' and M is a possible diagnosis.

That is, the normal physical model is selectively changed until it predicts (or is compatible with) the aberrant observations. Each change corresponds to an abnormality or malfunction. For example, in de Kleer and Williams' method for diagnosis, such a change consists of suspending the constraints of a component in the model, i.e., the outputs of a malfunctioning component are considered to be unconstrained. Each set of changes that accounts for the observations is considered a possible diagnosis. The process of diagnosis is a search for each such set of changes.

An example might clarify the intended role of physical models. Consider the device in Figure 1, by now a familiar example in the literature. Mult-1, Mult-2, and Mult-3 are three multipliers, with A-E as their inputs and X-Z as their outputs. Add-1 and Add-2 are adders with X-Z as their inputs. Add-1 and Add-2 are adders with X-Z as inputs and F and G as outputs. Given specific inputs to this de-




Figure 1: Multiplier and Adder Example

vice, and assuming that the device is functioning properly, the outputs can be predicted. If the actual outputs are improper, malfunctions can be attributed to one or more components by tracing through the connections and determining whether abnormal behavior would account for the outputs. For example, if output F is improper, then one possibility is that Mult-1 has a malfunction because if Mult-1 was no longer behaving like a multiplier, then F would be affected. Further testing with different inputs can verify or rule out this hypothesis about Mult-1 [6, 11].

Note that in this example, experiential knowledge is not mentioned, nor is it needed. The model of the device—the behavior of the components and the interconnections between the components—provides sufficient knowledge for diagnosis to be efficiently performed. If models with similar properties could be constructed for other domains, then the knowledge acquisition bottleneck could be avoided, and the conclusions of knowledge-based systems would be better justified.

Difficulties in Fulfilling the Promise

We have been studying diagnosis in the domain of human pathologic gait (walking disorders) [13, 18]. In this domain, the goal of diagnosis is to determine the muscular and skeletal causes of the patient's abnormal gait motions. Examples of malfunctions include reduced ranges of motion of joints (contractures), joint pain, muscle weakness, and muscle spasticity. Typically, patients will have a known underlying disorder such as cerebral palsy or arthritis, which gives rise to the joint and muscle malfunctions. Diagnosing joint and muscle malfunctions is done to help determine what kind of treatments (e.g., physical therapy, braces, surgery on joints and muscles) will best correct the patient's gait.

The data is primarily of the following types: history, physical exam, and motion data. The patient's history includes information about past and present diagnoses and treatments and demographic data. The physical examination provides data about the range of motion of joints and strength of muscles. The motion data include the gait motions of the patient obtained through a special camera system. This measures joint angles, gait velocity, stride length, etc. Also, EMG measurements are taken while the patient is walking, which provides data on muscle activity.

Since human walking is subject to the laws of Physics, just as any other physical activity, it appears that deep-knowledge-diagnosis would be appropriate. Unfortunately, there are several difficulties in doing diagnosis based on physical models.

1. *Construction.* Domain models with sufficient predictive and explanatory power need to be constructed before deep-knowledge-diagnosis can proceed. However, quantitative modeling of human gait is still a challenging research topic [12]. This is not just a problem in our domain. More often than not, simulation of complex mechanical devices and biological processes are *open research problems*.

2. *Ambiguity.* Even if a domain model can be constructed, there is a problem of obtaining sufficiently detailed data for the domain model. If a quantitative simulation is to be performed, then precise measurements of the initial state and input parameters need to be obtained. In many domains, this presents no difficulties. For example, electronic circuits are sufficiently constrained and well-understood so that carefully selected measurements can give the state of the device. Unfortunately in gait analysis, many internal gait parameters cannot be directly or even indirectly measured by current technology, e.g., EMG data is a best a qualitative measure of muscle forces [20]. Generally in medical domains, many internal parameters cannot be accurately measured without overly invasive actions.

One answer to this problem (and part of our own solution) is the use of qualitative physical models [1, 8, 10, 15]. Such models still provide explanatory power in spite of ambiguous data. However, qualitative models introduce their own sources of ambiguity. As a rule, qualitative simulation does not predict a single sequence of states, but produces several alternative state sequences. Additional information is required to disambiguate between them [7]. Also, qualitative simulation might produce state sequences that are spurious [16].

3. *Computational complexity.* Diagnosis is inherently computationally complex. The number of possible diagnoses is combinatorially large. If n different malfunctions can occur, then there are 2^n possible sets of malfunctions. If each malfunction can be caused in m different ways, then there are 2^{mn} possibilities. This large hypothesis space is not just an abstract possibility. In pathologic gait, patients with underlying neurological disorders typically have several malfunctions, some of which are "primary" (due to the underlying disorder), while others are attempted compensations.

Clearly, there is a need to modify the assumption that physical models for performing efficient and accurate diagnosis can be readily constructed in all domains. A more reasonable assumption is that physical models can perform specific diagnostic *subtasks*. Experiential knowledge acquired from human experts is still needed to help guide the search through the hypothesis space.

The Subtasks That Physical Models and Experiential Knowledge Are Good For

The next question is to identify the respective roles that physical models and experiential knowledge can play. Unfortunately, many factors are domain-dependent, and no "task model" is sufficiently developed to clearly answer this question (see [3, 5, 17] for what has been developed). Our own experience (the next best thing) is the following.

As mentioned earlier, deep knowledge is usually associated with physical models, which are intended to have pre-

dictive and explanatory power. A physical model of pathologic gait needs to represent at least the muscles, the limb segments, and the interactions among them. Given a particular situation (joint positions, limb and trunk momentums, muscle and ground forces), it ideally should predict changes in position and momentum. A major difficulty is obtaining accurate data on muscle forces. As a consequence, we will rely on a *qualitative* physical model, based on knowledge about the motions that muscles control and on relative strengths of muscles. For example, the action of a muscle on a joint might be described as "causes flexion" as opposed to a differential equation. Such models have weak predictive capabilities. In our domain, we will at best be able to explain how a motion could be caused by a combinations of factors, but the ambiguity concerning the exact amount of force associated with each factor precludes even qualitatively accurate predictions.

Another difficulty is using the physical model to search for diagnostic hypotheses. Because the interactions of the components (muscles and limb segments) are highly complex (unlike the device of Figure 1), a particular motion could be caused in a large number of ways, especially when combinations of malfunctions are considered. Also, the effect of any single malfunction can propagate throughout the rest of the system. A sprained ankle, for instance, affects the whole gait, not just the motion of the ankle. A physical model then might be able to suggest *local, individual* causes for a particular abnormal motion, but searching for all possible causes for each abnormal motion and generating composite hypotheses based on just this information will be computationally prohibitive.

Can experiential knowledge be used for the diagnostic reasoning that is difficult to do with physical models? As is typical in knowledge-based systems, experience can provide rules that associate abnormal observations with the malfunctions that typically cause them. Hence, experiential knowledge can give valuable clues concerning what malfunctions should be considered. Such knowledge, though, is not very good for determining whether a hypothesized malfunction accounts for the observations in a particular case, and is no good for considering combinations of interacting malfunctions.

Thus, we can usually use (qualitative) physical models to suggest some of the possible causes of an observation and to determine what observations a hypothesis accounts for (explanatory coverage). Experiential knowledge can associate hypotheses with observations and suggest which hypotheses appear more likely than others. In general, reasoning based on experiential knowledge is good for generating individual malfunctions that appear likely, while model-based reasoning is best for testing explanatory coverage of malfunctions and combinations thereof. Figure 2 summarizes these conclusions.

Integrating Model-Based Reasoning with Experiential Reasoning

These considerations lead us to the following proposal based on using hypothesis assembly [14] to integrate model-based

Experiential Reasoning	Model-Based Reasoning
generate: malfunctions associated with observations	generate: malfunctions that cause observations
test: likelihood of malfunctions	test: explanatory coverage of malfunctions
good for: generating individual malfunctions	good for: testing combinations of malfunctions

Figure 2: Experiential vs. Model-Based Reasoning

reasoning with experiential reasoning. Hypothesis assembly is a general technique for constructing and critiquing composite hypotheses. It requests information via the following domain-dependent functions:

1. a function that rates the "plausibility" of a hypothesis. The likelihood of malfunctions as determined by experiential reasoning can be used to rate plausibility. Explanatory coverage can be used to break ties (i.e., when the likelihoods of two malfunctions are too close for meaningful comparison).
2. a function that rates the importance of observations. In pathologic gait, the amount of the difference between abnormal and normal for a particular motion parameter determines its importance.
3. a function that determines what hypotheses can explain an observation. The qualitative physical model can be used to suggest hypotheses that explain an observation.
4. a function that determines what observations a (composite) hypothesis does and does not explain. The qualitative physical model can be used to determine what observations are explained by a combination of malfunctions.

The subtasks that experiential knowledge and physical models are good for fit into hypothesis assembly quite well.

Hypothesis assembly uses this information to construct a composite hypothesis with the following properties:

- The composite hypothesis explains as many observations as possible in comparison with similar composite hypotheses. That is, no local change (addition/deletion of some part to/from the composite hypothesis) improves explanatory coverage.
- Each hypothesis part within the composite hypothesis is as plausible as possible, viz. in comparison to other hypothesis parts explaining some particular observation.
- The composite hypothesis is parsimonious, i.e., no hypothesis in the composite hypothesis is superfluous.

(A hypothesis within a composite hypothesis is superfluous if it can be removed without loss of explanatory coverage.)

Hypothesis assembly also critiques this composite hypothesis in comparison to other composite hypotheses. Thus, one composite hypothesis is selected and its goodness in comparison to other hypotheses is determined.

Hypothesis assembly, however, is not guaranteed to find the "correct" hypotheses. Given the difficulties in deep-knowledge-diagnosis discussed earlier, no method can be expected to guarantee truth. Nor is hypothesis assembly guaranteed to produce the "best" hypotheses according to normative criteria such as "most probable hypothesis that accounts for all the observations." Such criteria are computationally intractable [2]. We conjecture that hypothesis assembly is the best that can be done within the constraints of imperfect physical models and computational tractability.

Conclusion

It has been proposed that diagnosis should be based on physical models of the domain. However, several factors make it unlikely that diagnosis can be just be based on physical models. These factors include constructing a sufficiently powerful physical model, obtaining sufficiently accurate observations, and performing diagnosis efficiently. Diagnosis in the domain of human pathologic gait illustrates these problems. Our proposal is to integrate qualitative physical models with experiential knowledge so that both sources of information will be efficiently and effectively utilized. In particular, they can be integrated using the technique of hypothesis assembly, which constructs a composite diagnostic hypotheses with several desirable properties: explanatory power, plausibility, and parsimony. We speculate that the reasoning architecture of our system is generally applicable to complex domains in which a less-than-perfect physical model and less-than-perfect experiential knowledge need to be combined to perform diagnosis.

Acknowledgments

This research has been supported by grants 82048-02 and H133E80017 from the National Institute on Disability and Rehabilitation Research.

References

- [1] T. Bylander. A critique of qualitative simulation from a consolidation viewpoint. *IEEE Trans. Systems, Man, and Cybernetics*, 18(2):252-263, 1988.
- [2] T. Bylander, D. Allemang, M. C. Tanner, and J. R. Josephson. *Some Results Concerning the Computational Complexity of Abduction*. Technical Report, Lab. for AI Research, CIS Dept., Ohio State U., Columbus, OH, 1988.
- [3] B. Chandrasekaran. Generic tasks in knowledge-based reasoning: High-level building blocks for expert system design. *IEEE Expert*, 1(3):23-30, 1986.
- [4] B. Chandrasekaran and S. Mittal. Deep versus compiled knowledge approaches to diagnostic problem-solving. *Int. J. Man-Machine Studies*, 19:425-436, 1983.
- [5] W. J. Clancey. Heuristic classification. *Artificial Intelligence*, 27(3):289-350, 1985.
- [6] R. Davis. Diagnostic reasoning based on structure and function. *Artificial Intelligence*, 24:347-410, 1984.
- [7] J. de Kleer and J. S. Brown. Assumptions and ambiguities in mechanistic mental models. In *Mental Models*, pages 155-190, Lawrence Erlbaum, Hillsdale, NJ, 1983.
- [8] J. de Kleer and J. S. Brown. A qualitative physics based on confluences. *Artificial Intelligence*, 24:7-83, 1984.
- [9] J. de Kleer and B. C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97-130, 1987.
- [10] K. D. Forbus. Qualitative process theory. *Artificial Intelligence*, 24:85-168, 1984.
- [11] M. R. Genesereth. The use of design descriptions in automated diagnosis. *Artificial Intelligence*, 24:411-436, 1984.
- [12] H. Hemami. Modelling, control, and simulation of human movement. *Critical Reviews in Biomedical Engineering*, 13(1):1-34, 1985.
- [13] V. T. Inman, H. J. Ralston, and F. Todd. *Human Walking*. Williams & Wilkins, Baltimore, 1981.
- [14] J. R. Josephson, B. Chandrasekaran, J. W. Smith, and M. C. Tanner. A mechanism for forming composite explanatory hypotheses. *IEEE Trans. Systems, Man and Cybernetics*, 17(3):445-454, 1987.
- [15] B. J. Kuipers. Commonsense reasoning about causality: Deriving behavior from structure. *Artificial Intelligence*, 24:169-203, 1984.
- [16] B. J. Kuipers. Qualitative simulation. *Artificial Intelligence*, 29(3):289-338, 1986.
- [17] J. McDermott. *Preliminary Steps Toward a Taxonomy of Problem Solving Methods*. Technical Report, Dept. of Computer Science, Carnegie-Mellon Univ., Pittsburgh, PA, 1988.
- [18] J. Perry. Normal and pathologic gait. In *Atlas of Orthotics*, chapter 4, C. V. Mosby, St. Louis, 1985.
- [19] R. Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57-95, 1987.
- [20] S. R. Simon. Kinesiology—its measurement and importance to rehabilitation. In V. L. Nickel, editor, *Orthopedic Rehabilitation*, chapter 5, Churchill Livingstone, New York, 1982.